

AMENDMENTS TO THE CLAIMS

1-4. (Cancelled)

5. (Currently Amended) The method according to claim 118, wherein the logical rules comprise a joining rule, and wherein the query plan comprises selecting a key responsively to the joining rule, and joining the data from two or more of the data sources using the key, and

wherein the selecting a the-key comprises analyzing the data so as to select one or more fields in the two or more of the data sources for use as the key so as to provide a desired statistical probability that the data will be joined correctly.

6-7. (Cancelled)

8. (Previously Presented) The method according to claim 118, wherein the logical rules comprise an access rule, and wherein the query plan comprises selecting at least one of the data sources for use in generating the response responsively to the access rule as applied to the user who submitted the query.

9-27. (Cancelled)

28. (Currently Amended) A method for data accessing data in a computer system including a set of diverse data sources comprising data, the method comprising:

defining an ontology for application to said a-set of diverse data sources comprising data;

defining data access rights with respect to the ontology; and

controlling user access to the data responsively to the ontology of the data and the access rights applicable thereto

collecting information in an ongoing manner regarding characteristics of the data sources;

storing said information in a storage device on said computer system;

employing at least one of wrappers and agents to automatically discover new or altered data sources associated with versions of said ontology;

automatically generating new ontology mappings of said data sources and said data access rights to said versions of said ontology in order to accommodate said new or altered data sources, without disrupting responses to queries on the basis of earlier versions of said ontology;

automatically generating new versions of said wrappers and agents to correspond with new versions of said ontology; and

automatically generating new versions of query plans to correspond with new versions of said ontology,

said automatically generating new ontology mappings, said automatically generating new versions of said wrappers and agents and said automatically generating new versions of query plans taking place without disruption of query response functionality of said computer system.

29. (Currently Amended) The method according to claim 28, wherein defining the ontology comprises specifying a user ontology, and wherein defining the data access rights comprises assigning a classification to a user according to the user ontology, and also comprising wherein controlling the user access comprises comparing the classification to the access rights applicable to the data.

30. (Original) The method according to claim 29, wherein the diverse data sources are distributed among a set of autonomous organizations comprising at least first and second organizations, and wherein assigning the classification comprises classifying the user according to an organizational affiliation of the user so as to control access by users in the first organization to the data sources held by the second organization.

31. (Currently Amended) The method according to claim 28, and also comprising wherein controlling the user access comprises receiving a query from a user to access the data in the data sources, determining a query plan for responding to the query by selecting one or more of the data sources responsively to the ontology such that the access rights permit the user to access the data in the one or more of the data sources, and generating a response to the query in accordance with the query plan.

32-33. (Cancelled)

34. (Currently Amended) The method according to claim 119, wherein the data sources are distributed among a set of autonomous organizations comprising at least first and second organizations, wherein the user submitting the query belongs to the first organization, and wherein determining the query plan comprises selecting, responsively to the performance characteristics, at least one of the data sources ~~resources~~ of the second organization for use in responding to the query.

35-105. (Cancelled)

106. (Currently Amended) A computer software product, comprising a computer-readable medium in which program instructions are stored, which instructions, when read by a computer, cause the computer to:

receive a definition of an ontology for application to a set of diverse data sources comprising data and a definition of data access rights with respect to the ontology; ~~and to control user access to the data responsively to the ontology of the data and the access rights applicable thereto~~

collect information in an ongoing manner regarding characteristics of the data sources;

store said information in a computer-readable manner;

employ at least one of wrappers and agents to automatically discover new or altered data sources associated with versions of said ontology;

automatically generate new ontology mappings of said data sources and said data access rights to said versions of said ontology in order to accommodate said new or altered data sources, without disrupting responses to queries on the basis of earlier versions of said ontology;

automatically generate new versions of said wrappers and agents to correspond with new versions of said ontology; and

automatically generate new versions of query plans to correspond with new versions of said ontology,

wherein said automatically generate new ontology mappings, said automatically generate new versions of said wrappers and agents and said automatically generate new versions of query plans take place without disruption of query response functionality of said product.

107. (Cancelled)

108. (Previously Presented) The product according to claim 106, wherein the diverse data sources are distributed among a set of autonomous organizations comprising at least first and second organizations, and wherein the instructions cause the computer to classify a user according to an organizational affiliation of the user so as to control access by users in the first organization to the data sources held by the second organization.

109-117. (Cancelled)

118. (Previously Presented) The method according to claim 28, and comprising:

associating with the ontology one or more logical rules applicable to semantics of the data in the data sources;

receiving a query from a user regarding the data;

determining a query plan for responding to the query by selecting one or more of the data sources responsively to the ontology and by identifying an operation to be applied to the data responsively to the applicable logical rules; and

generating a response to the query in accordance with the query plan.

119. (Previously Presented) The method according to claim 28, and comprising:

collecting information regarding a topology and performance characteristics of the data sources;

receiving a query from a user regarding the data;

determining a query plan responsively to the query and to the information regarding the

topology and performance characteristics; and

generating a response to the query in accordance with the query plan.

120-121. (Cancelled)

122. (Previously Presented) The product according to claim 106, wherein the instructions cause the computer to associate with the ontology one or more logical rules applicable to the semantics of the data in the data sources, and further cause the computer, upon receiving a query from a user regarding the data, to determine a query plan for responding to the query by selecting one or more of the data sources responsively to the ontology and by identifying an operation to be applied to the data responsively to the applicable logical rules, and to generate a response to the query in accordance with the query plan.

123. (Previously Presented) The product according to claim 106, wherein the instructions cause the computer to collect information regarding a topology and performance characteristics of the data sources, and further cause the computer, upon receiving a query from a user regarding the data, to determine a query plan responsively to the query and to the information regarding the topology and performance characteristics, and to generate a response to the query in accordance with the query plan.